

Requirements of an Integrated Formal Method for Intelligent Swarms

M. G. Hinchey
Information Systems Division
NASA GSFC
Greenbelt, MD, 20771, USA

Michael.G.Hinchey@nasa.gov

C. A. Rouff
SAIC
Advanced Concepts B.U.
McLean, VA, USA

rouffc@saic.com

J. L. Rash
W. F. Truszkowski
Information Systems Division
NASA GSFC
Greenbelt, MD, 20771, USA

ABSTRACT

The use of swarm technologies has become prevalent in a variety of application domains: medical, bioinformatics, military/defense, surveillance, even internet television broadcasting. Future NASA missions will exploit such technologies to enable spacecraft to be sent where heretofore it was impossible, to ensure greater protection of space assets, and to increase the likelihood of mission success. We describe some of the basic concepts of swarms, and discuss the requirements of a formal method suitable for use with swarm-based systems. We also present some findings of our FAST (Formal Approaches to Swarm Technologies) project, which is attempting to identify a suitable integrated formal method for this task.

Categories and Subject Descriptors

D.2.4 [Software/Program Verification]: Formal Methods; D.2.1 [Requirements/Specifications]: methodologies

General Terms

Design, Verification

Keywords

formal methods, verification, integrated/hybrid methods, swarms

1. INTRODUCTION

We are all familiar with swarms in nature. The mere mention of the word “swarm” conjures up images of large groupings of small insects, such as bees (apiidae) or locusts (acridiidae), each insect having a simple role, but with the swarm as a whole producing complex behavior [17].

Strictly speaking, such emergence of complex behavior is not limited to swarms, and we see similar complex social structures occurring with higher order animals and insects that don’t swarm *per se*: colonies of ants, flocks of birds, packs of wolves, etc. These groupings behave like swarms in many ways. With wolves, for

example, the elder male and female (alpha male and alpha female) are accepted as leaders who communicate with the pack via body language and facial expressions. Moreover, the alpha male marks the territory of the pack, and excludes wolves that are not members of the pack.

The idea that swarms can be used to solve complex problems has been taken up in several areas of computer science, which we will briefly introduce in Section 2. The term “swarm” in this paper refers to a large grouping of simple components working together to achieve some goal and produce significant results. The term should not be taken to imply that these components fly (or are airborne); they may equally well be on the surface of the Earth, under the surface, under water, or indeed operating on other planets.

We will describe NASA’s motivation for using swarms in future exploration missions. We will describe one particular mission, currently in the concept stage, and examine *why* this (and similar systems) *must* exhibit autonomic properties [17].

2. SWARMS AND INTELLIGENCE

Swarms consist of a large number of simple entities that have local interactions (including interactions with the environment) [2]. The result of the combination of simple behaviors (the microscopic behavior) is the emergence of complex behavior (the macroscopic behavior) and the ability to achieve significant results as a “team” [4].

Intelligent swarm technology is based on swarm technology. The individual members of the swarm also exhibit independent intelligence [3]. With intelligent swarms, members of the swarm may be heterogeneous or homogeneous. Even if members start as homogeneous, due to their differing environments, they may learn different things, develop different goals, and therefore become a heterogeneous swarm. Intelligent swarms may also be made up of heterogeneous elements from the outset, reflecting different capabilities as well as a possible social structure.

Agent swarms are being used as a computer modeling technique and have also been used as a tool to study complex systems [15]. Examples of simulations that have been undertaken include swarms of birds [6, 24], as well as business and economics [23], and ecological systems [30].

In *swarm simulations*, each of the agents is given certain parameters that it tries to maximize. In terms of bird swarms, each bird tries to find another bird to fly with, and then flies off to one side and slightly higher to reduce its drag. Eventually the birds form flocks. Other types of swarm simulations have been developed that exhibit unlikely emergent behavior. These emergent behaviors are the sums of often simple individual behaviors, but, when aggregated, form complex and often unexpected behaviors. Swarm behavior

is also being investigated for use in such applications as internet-based television broadcasting, telephone switching, network routing, data categorizing, and shortest path optimizations.

Swarm intelligence techniques (note the slight difference in terminology from “intelligent swarms”) are population-based stochastic methods used in combinatorial optimization problems, where the collective behavior of relatively simple individuals arises from their local interactions with their environment to give rise to the emergence of functional global patterns. Swarm intelligence represents a metaheuristic approach to solving a wide variety of problems.

Swarm robotics refers to the application of swarm intelligence techniques to the analysis of swarms where the embodiment of the “agents” is as physical robotic devices.

3. NASA SWARM TECHNOLOGIES

Future NASA missions will exploit new paradigms for space exploration, heavily focused on the (still) emerging technologies of autonomous and autonomic systems [35, 36]. Traditional mission concepts, reliant on one large spacecraft, are being complemented with mission concepts that involve several smaller spacecraft, operating in collaboration, analogous to swarms in nature. This offers several advantages: the ability to send spacecraft to explore regions of space where traditional craft simply would be impractical, greater redundancy (and, consequently, greater protection of assets), and reduced costs and risk, to name but a few. Planned missions entail the use of several unmanned autonomous vehicles (UAVs) flying approximately one meter above the surface of Mars, which will cover as much of the surface of Mars in three seconds as the now famous Mars rovers did in their entire time on the planet; the use of armies of tetrahedral walkers to explore the Martian and Lunar surface; constellations of satellites flying in formation; and the use of miniaturized pico-class spacecraft to explore the asteroid belt.

These new approaches to exploration missions simultaneously pose many challenges. The missions will be unmanned and necessarily highly autonomous. They will also exhibit the classic properties of autonomic systems, being self-protecting, self-healing, self-configuring, and self-optimizing. Many of these missions will be sent to parts of the solar system where manned missions are simply not possible, and to where the round-trip delay for communications to spacecraft exceeds 40 minutes, meaning that the decisions on responses to problems and undesirable situations must be made *in situ* rather than from ground control on Earth. The degree of autonomy that such missions will possess would require a prohibitive amount of testing in order to accomplish system verification. Furthermore, learning and adaptation towards continual improvements in performance will mean that emergent behavior patterns simply cannot be fully predicted through the use of traditional system development methods. The result is that formal specification techniques and formal verification will play vital roles in the future development of NASA space exploration missions.

3.1 ANTIS: A Concept Mission

Autonomous Nano Technology Swarm (ANTS) is a joint NASA Goddard Space Flight Center and NASA Langley Research Center collaboration to develop revolutionary mission architectures and exploit artificial intelligence techniques and paradigms in future space exploration. The mission will make use of swarm technologies for both spacecraft and surface-based rovers.

ANTS consists of a number of concept missions:

SARA: The Saturn Autonomous Ring Array will launch 1000 pico-

class spacecraft, organized as ten sub-swarms, each with specialized instruments, to perform *in situ* exploration of Saturn’s rings, by which to understand their constitution and how they were formed. The concept mission will require self-configuring structures for nuclear propulsion and control, which lies beyond the scope of this paper. Additionally, autonomous operation is necessary for both maneuvering around Saturn’s rings and collision avoidance.

PAM: Prospecting Asteroid Mission will also launch 1000 pico-class spacecraft, but here with the aim of exploring the asteroid belt and collecting data on particular asteroids of interest. PAM is described below in Section 3.1.1.

LARA: ANTS Application Lunar Base Activities will exploit new NASA-developed technologies in the field of miniaturized robotics, which may form the basis of remote landers to be launched to the moon from remote sites, and may exploit innovative techniques (described below in Section 3.1.2) to allow rovers to move in an amoeboid-like fashion over the moon’s uneven terrain.

Since SARA and PAM have many issues in common (as regards autonomous operation), we will concentrate on PAM in the following. Section 3.1.2 describes the unique technologies that are planned for the LARA (and other) concept missions.

3.1.1 PAM

The ANTS PAM (Prospecting Asteroid Mission) concept mission [10, 11, 35, 36] will involve the launch of a swarm of autonomous pico-class (approximately 1kg) spacecraft that will explore the asteroid belt for asteroids with certain characteristics.

Figure 1 gives an overview of the PAM mission concept [35]. In this mission, a transport ship, launched from Earth, will travel to a point in space where gravitational forces on small objects (such as pico-class spacecraft) are all but negligible. From this point, termed a Lagrangian, 1000 spacecraft, which will have been assembled *en route* from Earth, will be launched into the asteroid belt. As much as 60 to 70 percent of them are expected to be lost during the mission, primarily because of collisions with each other or with an asteroid during exploration operations, since, having only solar sails to provide thrust, their ability to maneuver will be severely limited. Because of their small size, each spacecraft will carry just one specialized instrument for collecting a specific type of data from asteroids in the belt. Approximately 80 percent of the spacecraft will be workers that will carry the specialized instruments (e.g., a magnetometer or an x-ray, gamma-ray, visible/IR, or neutral mass spectrometer) and will obtain specific types of data. Some will be coordinators (called leaders) that have rules that decide the types of asteroids and data the mission is interested in and that will coordinate the efforts of the workers. The third type of spacecraft are messengers that will coordinate communication between the rulers and workers, and communications with the Earth ground station.

The swarm will form sub-swarms under the control of a ruler, which contains models of the types of science that it wants to perform. The ruler will coordinate workers, each of which uses its individual instrument to collect data on specific asteroids and feed this information back to the ruler, who will determine which asteroids are worth examining further. If the data matches the profile of a type of asteroid that is of interest, an imaging spacecraft will be sent to the asteroid to ascertain the exact location and to create a rough model to be used by other spacecraft for maneuvering around the asteroid. Other teams of spacecraft will then coordinate to finish mapping the asteroid to form a complete model.




Figure 1: NASA's Autonomous Nano Technology Swarm (ANTS) mission scenario.

3.1.2 SMART

The ANTS SMART (Super Miniaturized Addressable Reconfigurable Technology) architectures were initiated at Goddard Space Flight Center (GSFC) to develop new kinds of structures capable of:

- goal-oriented robotic motion,
- changing form to optimize function (morphological capabilities),
- adapting to new environmental demands (learning and adaptation capabilities), and
- repairing-protecting itself (autonomic capabilities).

The basic unit of the structures is a tetrahedron (Figure 2) consisting of four addressable nodes interconnected with six struts that can be reversibly deployed or stowed. More complex structures are formed from interconnecting these reconfigurable tetrahedra, making structures that are scalable, and leading to massively parallel systems. These highly-integrated 3-dimensional meshes of actuators/nodes and structural elements hold the promise of providing a new approach to robust and effective robotic motion. The current working hypothesis is that the full functionality of such a complex system requires fully autonomous intelligent operations at each node.

The tetrahedron (tet) “walks” by extending certain struts, changing its center of mass and “falling” in the desired direction. As the tetrahedral structure “grows” by interfacing more and more tets, the falling motion evolves to a smoother walking capability, i.e., the smoother walking-climbing-avoiding capabilities emerge from the orchestration of the capabilities of the tetrahedra involved in the complex structure.

Currently, the basic structure, the tetrahedron, is being modeled as a communicating and cooperating/collaborating four-agent system with an agent associated with each node of the tetrahedron. An agent, in this context, is an intelligent autonomous process capable of bi-level deliberative and reactive behaviors with an intervening neural interconnection (the structure of the neural basis function [9]). The node agents also possess social and introspective behaviors. The problem to be solved is to scale this model up to




Figure 2: Basic unit of tetrahedral structures.

one capable of supporting autonomous operation for a 12-tet rover (a structure realized by the integration of 12 tets in a polyhedral structure). The overall objective is to achieve autonomous robotic motion of this structure. (See <http://ants.gsfc.nasa.gov> to view animations of the tetrahedon-based walking capabilities currently being modeled as multi-agent systems.)

3.2 Other NASA Swarm-Based Missions

An autonomous space exploration system is currently under development at Virginia Tech, funded by the NASA Institute for Advanced Concepts (NIAC).

The system consists of a swarm of low altitude, buoyancy-driven gliders for terrain exploration and sampling, a buoyant oscillating wing that absorbs wind energy, and a docking station that can be used to anchor the energy absorber, charge the gliders, and serve as a communications relay.

The work builds on success with underwater gliders currently used for oceanography research. The intent is to develop low-cost planetary exploration systems that can run autonomously for years in harsh environments such as in the sulfuric acid atmosphere of Venus, or on Titan (the largest of Saturn's moons).

3.3 NASA Constellations

We may consider constellations—several spacecraft flying together in formation—to be a special case of swarms.

The ST5 mission, for example, which is scheduled for Spring 2006, will launch three identical spacecraft that will fly in a “string of pearls” formation, utilizing a single uplink/downlink to earth. While a mission based on three spacecraft cannot be expected to perform highly distributed work as envisaged in ANTS, for example, it is certainly possible for the large spacecraft (or satellites) to be used in a constellation.

Indeed, this is the approach taken in the NASA Constellation-X mission. Constellation-X involves the use of a small number of telescopes (currently four), in formation and working together to give the equivalent of a single X-ray telescope for observing black holes and other X-ray sources with greater resolution than possible before [17].

4. OTHER APPLICATIONS OF SWARMS

The behavior of swarms of bees has been studied as part of the BioTracking project at Georgia Tech [1]. To expedite the understanding of the behavior of bees where large scale robust behavior emerges from the simple behavior of individuals, the project videotaped the behavior of bees over a period of time, using a computer vision system to analyze data on sequential movements that bees use to encode the location of supplies of food, etc. The intention is

that such models of bee behavior can be used to improve the organization of cooperating teams of simple robots capable of complex operations. A key point is that the robots need not have *a priori* knowledge of the environment, nor is there direct communication between robots in the teams.

Research at Penn State University has focused on the use of particle swarms for the development of quantitative structure activity relationships (QSAR) models used in the area of drug design [7]. The research created models using artificial neural networks and k-nearest neighbor and kernel regression. Binary and niching particle swarms were used to solve feature selection and feature weighting problems.

Particle swarms have influenced the field of computer animation also. Rather than scripting the path of each individual bird in a flock, the Boids project [24] elaborates a particle swarm with the simulated birds being the particles. The aggregate motion of the simulated flock is much like that in nature: it is the result of the dense interaction of the relatively simple behaviors of each of the (simulated) birds, where each bird chooses its own path.

Much success has been reported from the use of Ant Colony Optimization, a technique that studies the social behaviors of colonies of ants, and uses these behavior patterns as models for solving difficult combinatorial optimization problems [13]. The study of ants and their ability to find shortest paths has lead to ACO solutions to the traveling salesman problem, as well as network and internet optimizations [12, 13].

Work at University of California Berkeley is focusing on the use of networks of Unmanned Underwater Vehicles (UUVs). Each UUV has the same template information, containing plans, sub-plans, etc., and relies upon this and its own local situation map to make independent decisions, which will result in cooperation between all of the UUVs in the network. Experiments involving strategies for group pursuit will be conducted in a shallow water pool.

4.1 Verifying Swarms

As mission software becomes more complex, testing and error-finding also become more difficult. This is especially true of highly parallel processes and distributed computing, both being characteristic of swarm-based systems.

Race conditions in these systems can rarely be found by inputting sample data and checking whether the results are correct. These types of errors are time-based and only occur when processes send or receive data at particular times, or in a particular sequence, or after learning occurs. To find these errors through testing, the software processes involved have to be executed in all possible combinations of states (state space) that the processes could collectively be in. Because the state space is exponential in the number of states, it becomes untestable with a relatively small number of elements in the swarm. Traditionally, to get around the state-space explosion problem, testers have artificially reduced the number of states of the system and approximated the underlying software using models.

Formal methods are proven approaches for ensuring the correct operation of complex interacting systems. Once written, a formal specification can be used to prove properties of a system correct and check for particular types of errors (e.g., race conditions), and can be used as input to a model checker. Verifying emergent behavior is one area that, unfortunately, most formal methods have not addressed well.

The *FAST (Formal Approaches to Swarm Technologies)* project surveyed formal methods techniques to determine whether any may be suitable for verifying swarm-based systems and their emergent

behavior [28, 29]. The project found that there are a number of formal methods that support the specification of either concurrency or algorithms, but not both. Though there are a few formal methods that have been used to specify swarm-based systems, the project found only two formal approaches that were used to analyze the emergent behavior of swarms.

Weighted Synchronous Calculus of Communicating Systems (WSCCS), a process algebra, was used by Tofts to model social insects [34], and by Sumpter, et al., to analyze the non-linear aspects of social insects [33]. X-Machines have been used to model cell biology [19, 20], and with modifications, the X-Machines model has potential for specifying swarms. Simulation approaches are being investigated to determine emergent behavior [15]. However, these approaches do not predict emergent behavior from the model, but rather model the emergent behavior after the fact.

The project has defined an integrated formal method, which is appropriate for the development of swarm-based systems [26]. Future work will concentrate on the application of the method in order to demonstrate its usefulness, and on the development of appropriate support tools.

5. PROPERTIES OF AN EFFECTIVE INTEGRATED FORMAL METHOD FOR SWARM TECHNOLOGIES

An effective formal method must be able to predict the emergent behavior of 1000 agents as a swarm, as well as the behavior of the individual agent. Crucial to the mission will be the ability to modify operations autonomously to reflect the changing nature of the mission and the distance and low bandwidth communications back to Earth. For this, the formal specification will need to be able to track the goals of the mission as they change, and to modify the model of the universe as new data comes in. The formal specification will also need to allow for specification of the decision making process to aid in the determination of which instruments will be needed, at what location, with what goals, etc.

Once written, the formal specification to be developed must be able to be used to prove properties of the system correct (e.g., the underlying system will go from one state to another or not into a specific state), check for particular types of errors (e.g. race conditions), as well as be used as input to a model checker.

From this we can see that the formal method must be able to track the models of the leaders and it must allow for decisions to be made as to when the data collected has met the goals. The ANTS mission details are still being determined and are changing as more research is performed. Therefore, the formal method must be flexible enough to allow for efficient changes and re-prediction of emergent behavior.

Bearing all of this in mind, the following list summarizes the properties necessary for effective specification and emergent behavior prediction of the ANTS swarm and other swarms, and looks to the existing formal methods to provide some of the desired properties.

Process representation: Processes can be specified using the various manifestations of transition functions.

Reasoning: Other forms of possibly non-standard logics may need to be employed to allow for intelligent reasoning with uncertain and possibly conflicting information.

Choosing Action Alternatives: A means of expressing probabilities and frequencies of events (as in WSCCS) is most beneficial in choosing between different enabled actions. A modi-

fied version of the WSCCS ability may be used to supply an algebra for choosing between possible actions.

Asynchronous messaging: Asynchronous messaging will need to be supported, as this is the most common type of messaging in swarm applications. This is not a significant problem as most synchronous messaging is implemented via asynchronous “handshakes”. There are variants of CSP and other process algebras that support asynchronous messaging, either by having all processes be receptive (as in Receptive Process Theory), or through infinite buffering as in ACSP.

Message buffering: Message buffering may be needed due to the possibly asynchronous nature of messaging between members of the swarm. Several asynchronous variants of CSP achieve this through infinite buffering.

Concurrent agent states for each spacecraft: This requirement is well supported by available process algebras.

Communication protocols between agents: Available process algebras are highly effective in this area.

Adaptability to programming: Any formal specification languages that are developed will need to keep in mind the ease of converting the formal specification to program code and as input to model checkers.

Determining whether goals have been met: The goals of each of the spacecraft are constantly under review. We will need to be able to specify a method by which the spacecraft will know when the goals have been met. A modification to X-Machines may be able to solve this since the goals could be tracked using X-Machines (effectively finite state machines with memory).

Method for determining new goals: Once goals have been met, new goals must be formed. We need to be able to specify a method for forming these goals.

Model checking: Model checking will help to avoid semantic inconsistencies in the specifications. Notations employed will need to be suitable for use as input to a model checker.

Tracking Models: X-Machines have the ability to track the universe model in memory but need a more robust way to detail what the model is, how it is created, and how it is modified.

Associating agent actions with priorities and/or frequencies: A formal method deemed appropriate requires a means of expressing the probability of certain actions being enabled, and the frequency with which this will occur.

Predicting emergent behavior: Current approaches are not robust enough for the purpose of predicting individual and swarm emergent behavior and will need to be enhanced by greater use of Probability, Markov Chains, and/or Chaos Theory.

Table 1 illustrates part of the results of the survey for mainstream formal methods. Table 2 compares a number of the integrated or combination formal methods surveyed. Table 3 compares methods that, as reported in the literature, have been applied to modeling or specifying swarm-based systems (whether computer-based swarms, or real swarms in nature).

A significant issue for specifying (and verifying) swarm behavior is support for analysis of emergent behavior of swarms. The idea of emergence is well known from biology, economics, and other

scientific areas. It is also prominent in computer science and engineering, but the concept is not so well understood by computer scientists and engineers, although they encounter it regularly. Emergence refers to the fact that “the whole is often greater than the sum of its parts.” That is, when various components whose behavior is well understood are combined within a single environment, they often demonstrate behavior that is unexpected, and/or often cannot be foreseen or explained from the behavior of any individual component. The corollary of this is that making changes to components of a system of systems, or replacing a sub-system within a system of systems, may often have unforeseen, unexpected, and completely unexplained ramifications for both overall system behavior and the behavior of other subsystems.

Although the survey identified a few formal methods that have been used to specify swarm-based systems, initially only two formal approaches were found that had been used to analyze the emergent behavior of swarms, namely Weighted Synchronous Calculus of Communicating Systems (WSCCS) and Artificial Physics [31]. Since the survey was completed, two other approaches that may prove valuable in analyzing emergent behavior—CommUnity [14] and CSP2B [5]—have been brought to our attention, although we have not as yet identified their use with swarm technologies *per se*.

5.1 Experience specifying swarm behaviors

There has not been significant work on specifying swarm behavior. Interestingly, most of the work that has been reported in the literature has been related to specifying the behavior of swarms or colonies of insects, and has been performed by biologists with the assistance of computer scientists using modified formal methods. The following is a brief description of some specification techniques that have been used for specifying social, swarm, and emergent behavior:

- Weighted Synchronous Calculus of Communicating Systems (WSCCS), a process algebra, was used by Tofts to model social insects [34]. WSCCS was also used in conjunction with a dynamical systems approach for analyzing the non-linear aspects of social insects [33].
- X-Machines have been used to model cell biology [20, 19], and modifications, such as Communicating Stream X-Machines [21] also seem to have potential for specifying swarms.
- Dynamic Emergent System Modeling Language (DESMML) [22], a variant of UML, has been suggested for use in modeling emergent systems.
- Cellular automata [37] have been used to model systems that exhibit emergent behavior (e.g., land use).
- Artificial Physics [31], which uses physics-based modeling to gauge emergent behavior, have been used to ensure formation flying as well as other constraints on swarms.

Simulation approaches have also been investigated to determine emergent behavior, after which a modeling technique is used to model that behavior. Such approaches do not model emergent behavior *a priori*, instead only after the fact, and were not considered.

5.2 Evaluation of Specification Methods

Based on the results of the survey, four formal methods were selected to be used for sample specification of part of the ANTS mission. These methods were: the process algebras CSP [18, 16, 25] and WSCCS [34, 33], X-Machines [21], and Unity Logic [8].

Table 1: Comparison of candidate formal methods for intelligent swarms

Name	Concurrency Support	Algorithm Support	Tool Support	Formal Basis	Used in Agent-Based Specs.	Used in Swarm-Based Specs.
Artificial Physics	Yes	Yes	Yes	Yes	Yes	Yes (limited)
B	No	Yes	Yes	(Mathematical) Yes (Set Theory & Pred. Logic)	Yes	No
BDI Logic	Yes	No	Yes	Yes (Logic)	Yes	Yes (limited)
CSP	Yes	No	Yes	Yes (Algebraic)	Yes	No
Finite State Machines	No	Yes	Yes	Yes (Formal Lang.)	Yes	No
Game Theory	Yes	No	Yes	Yes (Mathematical)	Yes	Yes
I/O Automata	Yes	Yes	Yes	Yes (Formal Lang.)	Yes	No
KARO	Yes	No	Yes (limited)	Yes (Logic)	Yes	No
Mathematical Analysis	Yes	No	Yes	Yes (Mathematical)	Yes	Yes
Petri Nets	Yes	No	Yes	Yes	Yes	No
Pi Calculus	Yes	No	Yes	Yes (Algebraic)	Yes	No
Real Time Logic	Yes	No	Yes	Yes (Logic)	No	No
SCR	No	Yes	Yes	Yes (Formal Lang.)	No	No
Statecharts	Yes	No	Yes	No (Formal Lang.)	Yes	No
UML	Yes	Yes	Yes	No	Yes	No
X-Machines	No	Yes	No (limited)	Yes (Formal Lang.)	Yes	No
Z	No	Yes	Yes	Yes (Set Theory/ Pred. Calc.)	Yes	No

Table 2: Comparison of integrated formal methods

Name	Concurrency Support	Algorithm Support	Tool Support	Formal Basis	Used in Agent-Based Specs.	Used in Swarm-Based Specs.
Communicating X-Machines	Yes	Yes	No	Yes	Yes	Yes
CSP-OZ	Yes	Yes	No	Yes	Yes	No
Object-Z and Statecharts	Yes	Yes	No	Yes	Yes	No
Temporal B	Yes	Yes	No	Yes	Yes	No
Temporal Petri Nets	Yes	No	No	Yes	Yes	No
Timed Communicating Object Z	Yes	Yes	No	Yes	Yes	No
Timed CSP	Yes	No	Yes	Yes	Yes	No
ZCCS	Yes	Yes	No	Yes	Yes	No

These were used to describe a virtual experiment, described in section 4.3.1. CSP was chosen as a baseline specification method because the team has had significant experience and success [27, 28] in specifying agent-based systems with CSP. WSCCS and X-Machines were chosen because they have already been used for specifying emergent behavior by others, apparently with some success. Unity Logic was also chosen because it had been successfully used for specifying concurrent systems and was a logic-based specification, which offered a contrast to the other methods.

DESML, Cellular Automata, Artificial Physics, and simulation approaches were not used even though they had been used for specifying or evaluating emergent behavior. DESML, though very interesting, was not used because it had not been used or evaluated outside of the thesis it was developed under (though we may be revisiting it at a future time). Cellular Automata were not selected because they did not have any built in analysis properties for emergent behavior and because they have been primarily used for simulating emergent systems (as described in the previous section). Though not used for the specification, it too may be revisited to examine its strengths. Artificial physics, which is very promising, was not selected because of the newness of the approach and because of the translation that must be done between physics and software behavior. Lastly, simulation techniques were not used due to the fact that verification cannot be undertaken using simulation. This is because there could be emergent or other undesirable behaviors occurring that are not visible or do not become apparent during a simulation, but may be there nonetheless. A formal technique is designed to find exactly these kinds of errors.

6. AN INTEGRATED SWARM FORMAL METHOD

Integrating the above methods seems to be the best approach for verifying cooperating swarm-based systems. Integrating the memory and transition function aspects of X-Machines with the priority and probability aspects of WSCCS and other methods may produce a specification method that will allow all the necessary aspects for specifying emergent behavior in the ANTS mission and other swarm-based systems.

The merging of these formal methods is currently being performed. Figure 3 illustrates the proposed integrated formal method. The approach being taken is to use a conserving integration [32] of the methods. In this type of formal methods integration, the base formalisms of the methods are maintained and relationships between the formalisms are developed to reflect the new formal method. This approach will preserve the strength of the underlying methods and allow a seamless specification of the ANTS mission, and the development of support tools using existing semantics of the methods.

7. CONCLUSIONS

A brief overview of swarm technologies has been presented with emphasis on their relevance for potential future NASA missions. Swarm technologies hold promise for complex exploration and scientific observational missions that require capabilities that would be unavailable in missions designed around single spacecraft.

NASA is pursuing further development of formal methods techniques and tools that can be applied in the development of swarm-based systems, to help achieve confidence in their correctness.

Acknowledgements

This work has been supported by the NASA Office of Systems and Mission Assurance (OSMA) through its Software Assurance

Research Program (SARP) project, Formal Approaches to Swarm Technologies (FAST), administered by the NASA IV&V Facility, and by NASA Goddard Space Flight Center, Software Engineering Laboratory (Code 581).

We are grateful to Dr. George Hagerman (Virginia Tech) for information on Virginia Tech's Self Sustaining Planetary Exploration Concept; to Dr. Walter Cedeño (Penn State University, Great Valley) for information on Swarm Intelligence; and to Drs. Jonathan Sprinkle and Mike Eklund (University of California at Berkeley) for information on research into UAVs being conducted at that institution.

8. REFERENCES

- [1] T. Balch, A. Feldman, and Z. Khan. Automatic classification of insect behavior using computer vision and behavior recognition. In *Proc. Second International Workshop on the Mathematics and Algorithms of Social Insects*, Atlanta, Georgia, USA, 16–17 December 2003.
- [2] G. Beni. The concept of cellular robotics. In *Proc. 1988 IEEE International Symposium on Intelligent Control*, pages 57–62. IEEE Computer Society Press, Los Alamitos, Calif., 1988.
- [3] G. Beni and J. Want. Swarm intelligence. In *Proc. Seventh Annual Meeting of the Robotics Society of Japan*, pages 425–428, Tokyo, Japan, 1989. RSJ Press.
- [4] E. Bonabeau and G. Théraulaz. Swarm smarts. *Scientific American*, pages 72–79, March 2000.
- [5] M. J. Butler. *csp2B : A Practical Approach To Combining CSP and B*. Declarative Systems and Software Engineering Group, Department of Electronics and Computer Science, University of Southampton, February 1999.
- [6] S. Carlson. Artificial life: Boids of a feather flock together. *Scientific American*, November 2000.
- [7] W. Cedenro and D. K. Agrafiotis. Combining particle swarms and k-nearest neighbors for the development of quantitative structure-activity relationships. *Int. J. Comput. Res.*, 11(4):443–452, 2003.
- [8] K. M. Chandy and J. Misra. *Parallel Program Design: A Foundation*. Addison-Wesley Publishing Company, 1988.
- [9] S. Curtis, M. Rilee, W. Truszkowski, C. Cheung, and P. Clark. Neural basis function control of super micro autonomous reconfigurable technology (SMART) nano-systems. In *Proc. First AIAA Intelligent Systems Technical Conference*. AIAA, 20–22 September 2004.
- [10] S. A. Curtis, J. Mica, J. Nuth, G. Marr, M. L. Rilee, and M. K. Bhat. ANTS (Autonomous Nano-Technology Swarm): An artificial intelligence approach to asteroid belt resource exploration. In *Proc. Int'l Astronautical Federation, 51st Congress*, October 2000.
- [11] S. A. Curtis, W. F. Truszkowski, M. L. Rilee, and P. E. Clark. ANTS for the human exploration and development of space. In *Proc. IEEE Aerospace Conference*, Big Sky, Montana, USA, 9–16 March 2003.
- [12] M. Dorigo and L. M. Gambardella. Ant colonies for the traveling salesman problem. *BioSystems*, 43:73–81, 1997.
- [13] M. Dorigo and Stützle. *Ant Colony Optimization*. MIT Press, 2004.
- [14] J. L. Fiadeiro. *Categories for Software Engineering*. Springer-Verlag, London, 2004.
- [15] D. E. Hiebeler. The swarm simulation system and individual-based modeling. In *Proc. Decision Support 2001*:

Table 3: Comparison of formal methods used for swarm specifications

Name	Concurrency Support	Algorithm Support	Tool Support	Formal Basis	Emergent Behavior Analysis	Used in Swarm-Based Specs.
Cellular Automaton	Yes	Yes	Yes	Yes (FSM)	No	Yes
Com. X-Machines	Yes	Yes	No	Yes (Formal Lang.)	No	Yes
Unity Logic	Yes	No	Yes (limited)	Yes (Logic)	No	Yes
WSCCS	Yes	No	Some (Prob. Workbench)	Yes (Process Alg.)	Yes (Markov Chain)	Yes
Artificial Physics	Yes	Yes	Yes	Yes (Mathematical)	Yes	Yes (limited)

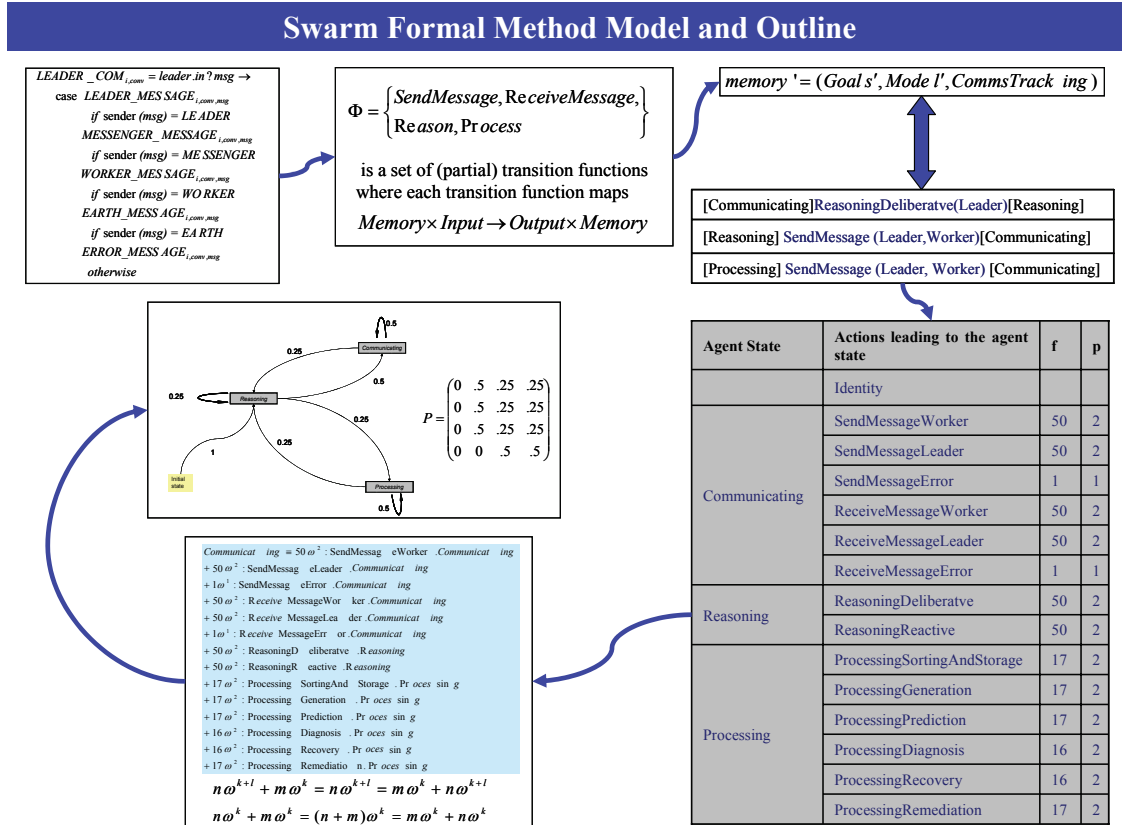


Figure 3: Proposed integrated formal method.

Advanced Technology for Natural Resource Management, Toronto, Canada, September 1994.

- [16] M. G. Hinchey and S. A. Jarvis. *Concurrent Systems: Formal Development in CSP*. International Series in Software Engineering. McGraw-Hill International, London, UK, 1995.
- [17] M. G. Hinchey, J. L. Rash, W. F. Truszkowski, C. A. Rouff, and R. Sterrit. Autonomous and autonomic swarms. In *Proc. The 2005 International Conference on Software Engineering Research and Practice (SERP'05)*, pages 36–42, Las Vegas, Nevada, USA, 27 June 2005. CSREA Press.
- [18] C. A. R. Hoare. *Communicating Sequential Processes*. Prentice Hall International Series in Computer Science. Prentice Hall International, Englewood Cliffs, NJ, 1985.
- [19] W. M. L. Holcombe. Mathematical models of cell biochemistry. Technical Report CS-86-4, Sheffield University, UK, 1986.
- [20] W. M. L. Holcombe. Towards a formal description of intracellular biochemical organization. Technical Report CS-86-1, Sheffield University, UK, 1986.
- [21] W. M. L. Holcombe. X-Machines as a basis for system specification. *Software Engineering*, 3(2):69–76, 1988.
- [22] J. R. Kiniry. The specification of dynamic distributed component systems. Master's thesis, California Institute of Technology, 1998.
- [23] F. Luna and B. Stefansson. *Economic Simulations in Swarm: Agent-Based Modelling and Object Oriented Programming*. Kluwer Academic Publishers, 2000.
- [24] C. W. Reynolds. Flocks, herds, and schools: A distributed behavioral model. *Computer Graphics*, 21(4):25–34, 1987.
- [25] M. L. Rilee, S. A. Boardsen, M. K. Bhat, and S. A. Curtis. Onboard science software enabling future space science and space weather missions. In *Proc. 2002 IEEE Aerospace Conference*, Big Sky, Montana, March 2002.
- [26] C. Rouff, A. Vanderbilt, M. Hinchey, W. Truszkowski, and J. Rash. Properties of a formal method for prediction of emergent behaviors in swarm-based systems. In *Proc. 2nd IEEE International Conference on Software Engineering and Formal Methods*, Beijing, China, September 2004.
- [27] C. A. Rouff, J. L. Rash, and M. G. Hinchey. Experience using formal methods for specifying a multi-agent system. In *Proc. Sixth IEEE International Conference on Engineering of Complex Computer Systems (ICECCS 2000)*, Tokyo, Japan, 2000. IEEE Computer Society Press, Los Alamitos, Calif.
- [28] C. A. Rouff, W. F. Truszkowski, M. G. Hinchey, and J. L. Rash. Verification of NASA emergent systems. In *Proc. 9th IEEE International Conference on Engineering of Complex Computer Systems*, Florence, Italy, April 2004. IEEE Computer Society Press, Los Alamitos, Calif.
- [29] C. A. Rouff, W. F. Truszkowski, J. L. Rash, and M. G. Hinchey. A survey of formal methods for intelligent swarms. Technical Report TM-2005-212779, NASA Goddard Space Flight Center, Greenbelt, Maryland, 2005.
- [30] M. Savage and M. Askenazi. Arborscapes: A swarm-based multi-agent ecological disturbance model. Working paper 98-06-056, Santa Fe Institute, Santa Fe, New Mexico, 1998.
- [31] W. M. Spears and D. F. Gordon. Using artificial physics to control agents. In *Proc. IEEE International Conference on Information, Intelligence, and Systems*, Charlotte, North Carolina, November 1999.
- [32] C. Suhl. RT-Z: An integration of Z and Timed CSP. In *Proc. 1st International Conference on Integrated Formal Methods (IFM99)*, York, United Kingdom, June 1999.
- [33] D. J. T. Sumpter, G. B. Blanchard, and D. S. Broomhead. Ants and agents: a process algebra approach to modelling ant colony behaviour. *Bulletin of Mathematical Biology*, 63(5):951–980, September 2001.
- [34] C. Tofts. Describing social insect behavior using process algebra. *Transactions on Social Computing Simulation*, pages 227–283, 1991.
- [35] W. Truszkowski, M. Hinchey, J. Rash, and C. Rouff. NASA's swarm missions: The challenge of building autonomous software. *IEEE IT Professional*, 6(5):47–52, September/October 2004.
- [36] W. F. Truszkowski, M. G. Hinchey, J. L. Rash, and C. A. Rouff. Autonomous and autonomic systems: A paradigm for future space exploration missions. *IEEE Transactions on Systems, Man and Cybernetics, Part C*, 2006 (to appear).
- [37] J. von Neumann. *Theory of Self-Reproducing Automata*. University of Illinois Press, Urbana, Illinois, 1996.